

Design Issues for the Solenoid Magnets for the Neutrino Factory Muon Cooling System

Michael A. Green
Lawrence Berkeley National Laboratory
Berkeley CA 94720

July 2000

The muon cooling system is the costliest part of the proposed neutrino factory. A major cost item for the cooling channel is the superconducting solenoids that surround the 200 MHz RF cavities for accelerating the muons after they have been cooled and the hydrogen absorbers that cool the muons going down the cooling channel. The size of the solenoids for the muon-cooling channel is set by the diameter of the 200 MHz RF cavities. This report discusses, in a general way, the steps needed to reduce the peak and average stresses in the solenoid and the solenoid cost. The design parameters that affect magnet training and quench protection are also discussed. This report presents the feasibility of building reduced size solenoids that will model the behavior of the full size magnets. The use of reduced scale model magnets is an attractive idea in that it would permit one to solve most of the fundamental superconducting magnet design problems at a reduced cost.

Introduction

The muon-cooling channel for the neutrino factory consists of several hundred meters of alternating liquid hydrogen absorbers to reduce both the longitudinal and transverse momentum of the muon and RF cavities to re-accelerate the muons in the longitudinal. In a perfect world the transverse momentum of the muons is removed while the longitudinal is preserved. The muons are carried in a solenoidal channel during the cooling and re-acceleration process. The on axis magnetic induction through most of the cooling channel is less than 5 T, but the solenoids must be large enough to accommodate the 200 MHz RF cavities. This means that the warm bore diameter of the solenoids around the cavities must be about 1.2 meters.

At least three types of cooling channels have been studied[1,2]. They are: 1) The first type of channel is the FOFO channel that contains alternating polarity solenoids with the polarity of the solenoids changing (the on axis field is a gradient field) at the RF cavity. The field in the region of the liquid hydrogen absorber is solenoidal rather than a gradient field. The assumed solenoids for the FOFO channel have all the same inside warm bore diameter with a peak on-axis induction of about 3.5 T. There are periodic gaps in the solenoid to provide RF to the cavities and refrigeration for the liquid hydrogen absorber. 2) The second type of channel is a flip channel. In this channel the solenoid all have the same polarity for long distances. In one or two places along the channel, the polarity of the solenoidal induction changes. The average induction along this channel is about 4 T. The continuous solenoids have periodic gaps between coils so that RF cavities can be supplied with RF and refrigeration can be supplied to the liquid hydrogen absorbers. The large magnetic forces associated with flipping the polarity of the field are confined to one or two places along the cooling channel. 3) The third type of channel the Super FOFO channel, which is a variation of the FOFO channel. The field flips periodically along the channel, but the gradient section occurs in the liquid hydrogen absorber rather than in the RF cavities. Since the liquid hydrogen absorbers have a smaller outside

diameter than the RF cavities, the solenoids where the change in field polarity occurs can be made smaller in diameter. The peak field along the channel in the standard super FOFO case is about 2.8 T. Increasing the field as one moves down the channel can potentially improve the level of muon cooling in the channel. Improved cooling is needed so that the cost of the re-circulating acceleration rings and the muon storage ring can be reduced. A question that one asks is how far can one go in increasing the field in the channel and how much will it cost compared to the standard case? The three cooling channels are illustrated in Figure 1 below. There are going to be many other muon-cooling schemes before a design is frozen.




Figure 1. Schematic Representations of Various Muon Cooling System Structures

The magnets shown in Figure 1 on the previous page have a number of things in common. There is at least one point along the channel where is a rapid change of the field or flux reversal. The magnets can not be continuous over the length of the channel. RF feed to the cavities must be made periodically along the length of the structure. There must also be cryogenic and gas services to the liquid hydrogen absorbers. Except for the flip polarity structure, the shape of the magnetic field is important to the beam dynamics. The periodicity of the magnetic field can contribute to beam loss as the muons travel down the structure. In all three cases, the magnets are large enough to be stress and strain limited. The cost of the magnet is dominated by the mass in the coil needed to hold the magnetic forces. Quench protection is another issue that is common to the three cooling channels shown in Figure 1.

This report looks at the muon cooling channel magnets. This report shows how the radial position of the solenoid coils affect the quality of the field in a straight channel and the gradient produced in the field flip regions. This report discusses what one can do to reduce the number of ampere-turns needed to generate a given field profile on axis. The average stress in the coils is discussed along with the peak stress condition (the so-called BJR criteria)[3]. This report directly relates quench protection to stored magnetic energy and stress. Finally this report will give the criteria for building reduced scale model magnets that can be used to solve real stress and quench protection problems before the full sized channel magnets are built.

Magnetic Field and Stress in Simple Solenoids

The simplest type of solenoid is the infinite solenoid. In infinite solenoid generate a uniform field across its bore. The central induction and the peak induction in the coil are the same. A solenoidal thin sheet of current can represent the infinite solenoid, but its properties are the same for a current carrying coils of finite thickness. The induction within the solenoid coil is uniform determined by the current per unit length along the axis. The induction outside the coil is zero. The induction B in an infinite solenoid can be calculated using the following expression [4]:

$$B_z = \mu_0 nI \quad -1-$$

where B_z is the central induction in the solenoid; nI is the number of ampere turns per meter; and μ_0 is the permeability of air. The average current density in the windings J is uniform along the length of the solenoid. For a solenoid coil of thickness t the average current density in the windings is

$$J = \frac{nI}{t} = \frac{B_z}{\mu_0 t} \quad -2-$$

Since a solenoid behaves like a magnetic pressure vessel, the average stress in the coil can be calculated using the pressure vessel equation as given as follows [5,6]:

$$\sigma = \frac{P_m R}{t} = \frac{B_z^2}{2\mu_0} \frac{R}{t} \quad -3-$$

where P_m is the magnetic pressure defined as $B_z^2 / 2\mu_0$; R is the coil radius and t is coil thickness. An alternative way of representing the average stress in the coil is to define the coil Lorentz force as the product of current in the coil nI times the average induction in the coil $B_z/2$. The alternative expression for the average stress in the coil takes the following form:

$$\sigma = \frac{nIB_z}{2} \frac{R}{t} = \frac{B_z^2}{2\mu_0} \frac{R}{t}, \quad -4-$$

which is the same as given by Equation 3.

The peak induction in an infinite solenoid is the central induction B . If the coils were divided into very small sheets of current not connect mechanically with the other current sheets, the peak stress in this sheet σ_p would be the induction in the innermost sheet B_z times the radius of the sheet R_i times the current density in the sheet J . This is the so-called BRJ design criterion. For an infinite solenoid the maximum possible peak stress in the coil σ_p is:

$$\sigma_p = B_z R J = \frac{R B_z n I}{t} = \frac{B_z^2 R_i}{\mu_0 t} \quad -5-$$

When the infinite solenoid coil is thin so that $R = R_i$ the peak possible stress on the inside of the coil is twice the average stress in a thin solenoid coil. Since the conductors in a real solenoid are tied together mechanically, the peak and average stresses in the coil are much closer to each other than a factor of two.

Before proceeding further it is useful to comment on the so-called BRJ design criteria. John Miller of the National High Magnetic Field Laboratory and others have stated that peak value of BRJ in a solenoid coil should not exceed about 350 MPa [3]. It should be pointed out that this is arbitrary design criteria that only applies to solenoids. B in this case is defined as the peak value of the magnetic induction in a direction parallel to the solenoid axis in the coil. J is defined as the coil current density; but more correctly it should be the current density in the parts of the coil that carry current and hoop stress. R is the radius of the location in the coil where the peak induction parallel to the axis is located. In thick solenoids where t is greater than $0.1 R$, this criterion becomes of some importance. When solenoidal coils are short and thick, the BRJ criteria can have an effect on the local strain in the coil, which in turn is reflected in coil motion. The $BRJ = 350$ MPa is only a design guide. Designing coil with a BRJ greater than about 450 MPa should be done with great caution. One should know where the coil is carrying the forces and one should understand what the relative motion between one part of the coil and another is. The relative motion between coil parts is part of the cause of magnet training. For most solenoids the peak value of BRJ will be more than twice the average stress in the coil. For many of the magnets considered for the muon collider cooling system the peak BRJ value is more than four times the average stress in the coil.

It is useful to look at the relationship between coil stress and the magnet stored-energy in a thin infinite solenoid. The stored energy per unit length for a thin infinite solenoid can be calculated using the following expression:

$$E = \frac{B_z^2}{2\mu_0} \pi R^2 \quad -6-$$

where E is the stored energy of the solenoid per unit length; R is the radius of the thin solenoid; and all other coefficients in the equation have been previously defined.

By comparing Equations 3 and 6 one can see that there is a direct correlation between the stored energy per unit length of the solenoid and the average stress in the solenoid winding. The solenoid average stress σ stated in terms of the solenoid stored energy per unit length E , the coil radius R and the coil thickness t can be calculated using the following expression:

$$\sigma = \frac{E}{\pi R t} \quad -7-$$

This says that for a given coil radius R , a given average design stress σ , the relationship between coil thickness t and coil stored energy E is as follows:

$$t = \frac{E}{\pi R \sigma} \quad -8-$$

This also suggests that the coil thickness t must go up as B^2 for a given design stress σ and coil radius R .

Scaling Laws for Large Solenoids Based on Stress Limits

Building a scale model of a cooling module solenoid system has been suggested, because the cost of the scale model is a fraction of the cost of a full-scale cooling solenoid magnet system. For example, the fabrication cost of a half scale solenoid may be a fifth to a quarter of the fabrication cost of a full-scale solenoid magnet system. Engineering done on the small-scale prototype is nearly that same as for the full-size magnet. The question that comes to mind is can a smaller scale solenoid be built address the stress and strain issues associated with the full-scale solenoid? The answer is yes, provided the scale model solenoid is not too small. From a stress and strain standpoint, scaling is possible within certain limits. Equations 3 and 4 suggests that there are a set of scaling laws that can be applied to solenoids from the standpoint of average stress and peak stress in the winding. When one designs a solenoid, one designs the magnet with a particular average value of stress (strain) in mind. If the coils are correctly scaled, the peak stresses should also scale within a few percent.

Equation 3 suggests a scaling law that can be used for a large solenoid. From equation 3, the following scaling law for coil thickness as a function of coil radius emerges:

$$t = \frac{B_z^2}{2\mu_0 \sigma} R \quad -9-$$

Equation 9 suggests that for a given average design stress σ in the coils and the coil support structure, and a given magnetic induction in the bore B_z , that coil thickness t is proportional to R and proportional to B_z squared.

The current density in the winding as a function of the average stress in the coil σ , the central magnetic induction B , and the coil radius R can be calculated using the following expression:

$$J = \frac{\sigma}{B R} \quad -10-$$

In order to achieve the same induction as the full-scale magnet, the current density in the conductor must scale inversely with R , because the number of ampere-turns per unit length does not change, but the coil thickness goes down with R . For a given design stress in the coil, the current density in the winding must go down with central induction, because the coil thickness must go up with the induction squared (see Equation 9). In general, the scaling equations 9 and 10 should be applicable to other solenoid configurations besides simple infinite solenoids.

There are two factors that limit the applicability of scaling. The first is the fraction of the coil package that is superconductor. The second is the fraction of the coil package that is electrical insulation. The thermal insulation space in the cryostat can not easily be scaled without affecting cryostat performance, but oddly enough the thickness of critical elements (such as the vacuum vessel) of the cryostat do scale with the coil radius. Scaling of the cryostat elements are discussed in Reference 7.

First lets look at the superconductor. In general, a large solenoid superconducting coil will have a conductor that is a small fraction superconductor (say 10 parts aluminum or copper to 1 part superconductor). A half scale model must have twice as much superconductor as compared to the normal metal in the matrix. A 9 to 1 matrix to superconductor ratio conductor in the full-scale magnet will translate out to a 4 to 1 matrix to superconductor ratio for a half-scale model.

If we are talking about niobium titanium as the superconductor, this is not a great problem. Niobium-titanium has roughly the same elastic modulus as copper (145 GPa). The yield stress for the Nb-Ti at 4 K is about 1200 MPa (173 ksi), whereas the yield stress for the copper in the conductor can be from a 200 to 400 MPa (29 to 58 ksi) depending how much the wire has been drawn in the fabrication process. One of the arguments behind the BRJ limit for coil design of 350 MPa is the possibility of yielding the matrix material in a coil. Niobium-tin is another issue. Niobium-tin conductor becomes almost dead soft during the heat treatment process. Increasing the fraction of niobium-tin in the conductor will affect the ability of the conductor to carry stress. Even with niobium-tin conductor, one can scale the magnet but one should not reduce the size of the magnet over a factor of 2.5. Full-scale magnet designs that call for a small copper to superconductor ratio conductor (say less than 3 to 1 copper to superconductor ratio) are not easily scaled because there can not be enough superconductor to carry the current in the coil in a stable way. Figure 2 below shows the J_c of Nb-Ti and Nb₃Sn as a function of B and T[8]. Figure 3 on the next page shows maximum coil J versus B and T for Nb-Ti and Nb₃Sn Coils.




Figure 2. Superconductor J_c for Nb-Ti and Nb₃Sn versus Induction at 4.2 K and 1.8 K




Figure 3. A Maximum Design Coil Current Density for Small Non-thin Nb-Ti And Nb₃Sn Coils as a Function of Induction at 4.2 K and 1.8 K

Coil insulation is the other factor that is affected by scaling. For the most part, the turn to turn, the layer to layer and ground plane insulation thickness can not be changed as the size of the magnet is reduced. For a coated conductor, the thickness of the coating is about 0.05 mm regardless of the size of the conductor. The layer to layer insulation thickness for a potted coil design is always 0.2 to 0.3 mm. The ground plane insulation can be scaled a factor of two say from 1.6 mm to 0.8 mm provided the ground plane insulation consists of overlapping layers of fiber-glass epoxy resin type material. The total thickness of the insulation in the coil relative to the radial and longitudinal dimensions of the coil affects the coil overall elastic modulus in the both directions. The reduced scale coils will have a lower overall modulus in both the radial and longitudinal directions. The reason for the reduced coil modulus is the fact that insulation has an elastic modulus that is 10 to 20 percent of that of the conductor. The reduced coil modulus has the effect of increasing the peak stress in the winding when the Lorentz force is applied. The pre-stress in the coil after cool-down is also reduced in the reduced scale coils. As a result, the reduced scale coil will be more prone to training than will be the full-scale coil. If one can make the reduced scale magnet system operate without training or excessive mechanical deflection, the full-scale magnet system is not likely to suffer from these defects.

Neither the reduced matrix to superconductor ratio nor the increased relative insulation thickness should prevent the scale model coil from representing the full-scale coil, provided the test coil is not smaller than say 40 percent scale. In either case, if the scale model magnet performs well the full-scale magnet should perform better. If scaling is done correctly, the field quality in the scale model magnet should be similar to the field quality in the full-scale magnet. The substantially reduced cost of the reduced scale coil makes the construction of a scale model muon cooling magnet module very attractive. From a stress standpoint the reduced scale magnet is the right thing to build.

Quench Propagation and Quench Protection in a Simple Solenoid

A second design condition that one must consider in solenoids is quench protection. As the coil quenches and becomes resistive, the stored energy of the magnet is deposited in the coil. When a magnet quenches, one worries about the temperature at the hot spot, which is usually located at the point where the normal region propagation in the coil begins. The effective hot spot temperature is limited by the organic insulation in the coil. It is not unreasonable to design a coil and its quench protection system for a hot spot temperature $T_M = 300$ K (27 C). The hot spot temperature of the coil depends on the resistivity of the matrix material, the volume specific heat for the matrix and superconducting material and the integrated current density in the conductor with time as the magnet quenches. An expression for the hot spot behavior of a quenching superconducting magnet is given as follows [9,10]:

$$F^* = \int_{T=4K}^{T_M} \frac{C(T)}{\rho(T)} dT = \frac{r+1}{r} \int_{t=0} J^2(t) dt \quad -11-$$

where F^* is the integrated current density function; $C(T)$ is the volume specific heat as a function of temperature T ; and $\rho(T)$ is the matrix material resistivity as a function of temperature. $J(t)$ is the current density in the superconductor plus matrix as a function of time t ; and r is the ratio of matrix area to superconductor area in the coil conductor.

Equation 11 directly relates the physical properties of the matrix material with the current density integral. The $r+1/r$ function takes into account the fact that when the superconductor goes normal only the matrix material can carry the current yet the superconductor is counted in the volume specific heat for the material. Table 1 below gives values for F^* as a function of the hot spot temperature T_M for copper and aluminum of various residual resistivity ratios (RRR defined as the ratio of the resistivity at 273 K over the resistivity at 4 K).

Table 1. The value of the Current Density Integral Function F^* as a function of Material, Hot Spot Temperature T_M , and the Residual Resistivity Ratio RRR

Material	$T_M = 100$ K	$T_M = 200$ K	$T_M = 300$ K	$T_M = 400$ K
Copper RRR=10	4.5×10^{16}	8.3×10^{16}	10.1×10^{16}	12.7×10^{16}
Copper RRR=30	5.3×10^{16}	9.1×10^{16}	11.9×10^{16}	14.5×10^{16}
Copper RRR=100	7.9×10^{16}	11.7×10^{16}	14.5×10^{16}	17.0×10^{16}
Copper RRR=200	8.8×10^{16}	12.6×10^{16}	15.4×10^{16}	17.9×10^{16}
Aluminum RRR=25	1.3×10^{16}	2.1×10^{16}	3.6×10^{16}	4.7×10^{16}
Aluminum RRR=100	2.2×10^{16}	3.4×10^{16}	4.9×10^{16}	6.0×10^{16}
Aluminum RRR=300	3.0×10^{16}	4.2×10^{16}	5.7×10^{16}	6.8×10^{16}
Aluminum RRR=1000	3.8×10^{16}	5.0×10^{16}	6.5×10^{16}	7.6×10^{16}
Aluminum RRR=3000	4.6×10^{16}	5.8×10^{16}	7.3×10^{16}	8.4×10^{16}

From Table 1 above it is clear that an aluminum matrix conductor has a lower current density integral for a given temperature than does a copper based conductor. The reason for this is the lower value of the volume specific heat and the higher value of the resistivity for a given RRR. (The 273 K value of resistivity for aluminum is about 1.6 times the 273 K value for the resistivity for pure copper.) It is hard to get a copper in superconductor that has an RRR that is much greater than 200. Pure aluminum is another matter, because RRR above 2000 are almost routine for a very pure aluminum

matrix in a conductor. From Table 1 one might think that it is not a good idea to use aluminum conductor in a superconducting magnet. For large relatively low field solenoids, the opposite is true. A solenoid with an aluminum based conductor has a matrix current density that is 1.6 to 2 times lower than for a typical copper based conductor magnet. The lower current density in the conductor means that the aluminum matrix conductor is better protected because the current decay time constants in the magnet are longer.

From Equation 11 one can determine the current decay time constant for a quenching magnet in order that the magnet hot spot temperature not be above T_M . An expression for the required exponential current decay time constant, for a magnet to be protected is given as follows:

$$\tau = \frac{2F^*(T_m)}{J_o^2} \frac{r}{r+1} \quad -12-$$

where F^* is defined by equation 11 and T_m is the maximum allowable hot spot temperature for the magnet. J_o is the current density in the superconductor at the start of the quench. (For many superconducting magnets, $J_o = 1.2$ to 1.4 times the average coil current density at full design current, because of the insulation in the windings.) The symbol r is the matrix area to superconductor area ratio for the conductor.

The current decay time constant calculated using Equation 12 applies to current decay due to the growing resistance of a quenching magnet or the current decay due to discharging the magnet through an external resistor. As long as the current decay time constant for the magnet during a quench is less than calculated using Equation 12, the magnet will have a hot spot temperature at the origin of the quench that is less than T_M . When one chooses to use an external resistor R_{ex} to protect a quenching magnet with a self-inductance L_1 , the value of this resistance is given by the following expression:

$$R_{ex} = \frac{r+1}{r} \frac{J_o^2}{2F^*(T_M)} L_1 \quad -13-$$

where L_1 is the self-inductance of the magnet or the protected section of the magnet. J_o , r and F^* are previous defined. The self-inductance of the magnet L_1 can be calculated directly or it can be calculated from the stored energy of the protected section E_o using the following expression:

$$L_1 = \frac{2E_o}{I^2} \quad -14-$$

where E_o is the stored energy of the section when the coil carries its design current I .

Once one has defined the coil design current, one can design the coil insulation system for the maximum allowable voltage during a quench V . One might argue that one should make V as large as possible. When one makes the quench voltages very large, one has to add more insulation to the coil, which reduces the packing fraction for the coil. For most practical potted coil designs, the quench design voltage can be set to be 1000 to 1500 volts. This brings us to the EJ^2 limit for a magnet, which is defined as follows [10,11]:

$$E_o J_o^2 = VIF^*(T_M) \frac{r+1}{r} \quad -15-$$

where E_o is the stored energy at the magnet design current I ; J_o is the current density in the superconductor plus matrix at the magnet design current; V is the discharge voltage across a quench

protection resistor, $F^*(T_M)$ is the current density integral function with time (also the $C/$ integral function with temperature) and r is the ratio of the matrix area to the superconductor area in the conductor.

Equation 15 ties magnet stored energy and superconductor plus matrix current density to the quench protection of the magnet. For a typical magnet, with a design current $I = 1000A$, protected by an external resistor that puts a voltage $V = 1000 V$ across the leads, the EJ^2 limit will be about $10^{23} A^2 m^{-4} J$ for a hot spot temperature T_M of 300 K. If the magnet is wound with an aluminum matrix superconductor the EJ^2 limit will be lower. In order for a magnet to operate at EJ^2 limits that are substantially higher than $10^{23} A^2 m^{-4} J$, the coil current must be increased or the magnet coil must be driven normal using external heaters or through quench back from a closely coupled secondary circuit.

The EJ^2 limit can also be tied to stress in the winding for a simple solenoid of infinite length. For a one-meter section in the middle of an infinite solenoid of radius R and thickness t , the EJ^2 limit takes the following form:

$$EJ^2 = \pi R^2 \frac{B_z^2}{2\mu_0} \frac{nI}{t} \quad -16-$$

where nI is the ampere turns per unit length; B is the central induction of the solenoid and R , t , μ_0 are previously defined. Equation 16 can be manipulated to put EJ^2 per unit length in terms of stress in the winding (assuming that only the conductor that carries current carries the stress. The form of the EJ^2 limit per unit length that includes stress is:

$$EJ^2 = \frac{2\pi}{\mu_0} \sigma^2 \quad -17-$$

where σ is the hoop stress in the conductor. If one applies the $BRJ = 350$ MPa rule to the infinite solenoid, the average stress $\sigma = 175$ MPa. At an average stress of 175 MPa, the EJ^2 limit is about $1.53 \times 10^{23} A^2 m^{-4} J$, which is very close to the EJ^2 limit imposed by quench protection. In order to increase the EJ^2 limit imposed by stress, one must off load the magnetic forces to a closely coupled member that can carry the stress. In large closely coupled thin solenoid magnets, the stress-carrying member is the same member that is coupled inductively to the coil, which the coil quench-backs from.

Small superconducting magnets do not need a quench-protection system. Magnets that can quench safely without a quench protection system go normal fast enough so that the resistance of the coil equals the value of R_{ex} calculated by Equation 13 in a time that is less than the value of t_q calculated by Equation 12. The increase in the resistance of a two dimensional magnet that is thin compared to its radius (say $t < 0.1R$) during a quench takes the following form [12]:

$$R(t) = \alpha R_{o2} v_L^2 t^2 \quad -18-$$

where $R(t)$ the coil resistance as a function of time t ; R_{o2} is the resistance coefficient for the magnet; v_L is the quench propagation velocity along the wire and α is the ratio of the turn to turn quench propagation velocity to the propagation velocity along the wire. The ratio of turn to turn quench velocity to quench velocity along the wire is measure of the thermal conductivity ratio in the two directions. The value of α can be estimated using the following expression:

$$\alpha = \frac{\rho k_i}{LT_c} \frac{a}{c} \frac{r+1}{r}^{0.5}$$

-19-

where ρ is the resistivity of the superconductor matrix material, k_i is the thermal conductivity of the insulation; L is the Lorenz number ($L=2.45 \times 10^{-8} \text{ W K}^{-1}$); T_c is the critical temperature of the superconductor; c is the turn to turn insulation thickness; a is the conductor width; and r is the area of the matrix over the area of the superconductor. For a typical coil where $\rho = 1.5 \times 10^{-10} \text{ m}$; $k_i = 0.1 \text{ W m}^{-1} \text{ K}^{-1}$; $T_c = 10 \text{ K}$; $a/c=10$ and $r = 5$, the value of $\alpha = 0.028$. For a one meter diameter coil, the quench will propagate about 0.088 m along the coil (in both directions) in the time that the normal region propagates half way around the coil, at which point the quench becomes one dimensional. To give one a feeling for the value of the propagation velocity along the wire, see measurements of quench propagation velocity in Figure 4 below.

Matrix Plus S/C Current Density (A m^{-2})




Figure 4. Measurements of Quench Velocity as a Function of Current Density in the Magnet Conductor for Three LBNL Potted Thin Solenoids

Figure 4 above shows a straight line that has been fitted to the measured data for the C coil, which is a two-meter diameter thin solenoid. An approximate equation for the fitting line is given by the following expression [10,13]:

$$v_L = 2 \times 10^{-14} J^{1.7}$$

-20-

where J is the current density in the conductor (superconductor plus matrix).

Studies of normal region propagation along the wire show that there is almost no dependence on r for either a copper or aluminum matrix conductor. Equation 20 is reasonably valid for a copper based

conductor from $J = 10^8 \text{ A m}^{-2}$ to above $J = 8 \times 10^8 \text{ A m}^{-2}$. For a conductor that has an aluminum matrix, the propagation velocity along the wire increases over a factor of three [14]. For wires with an ultra pure aluminum matrix, the quench propagation velocities along the wire are even larger [15]. A copper matrix niobium-tin conductor will have quench propagation velocities along the wire that are a factor of two to three lower. Quench propagation velocity along the wire does increase with magnetic field because less heat is needed to drive the superconductor normal.

Let's look at a one-meter long section of coil that is part of an infinite solenoid. This section has its own power supply and leads. If $R = 0.5 \text{ m}$, the value of J that will not exceed the BJR criteria of 350 MPa will be $1.4 \times 10^8 \text{ A m}^{-2}$. For a one-meter long section, the stored energy will be $7.8 \times 10^6 \text{ J}$. Thus the EJ limit will be $1.53 \times 10^{23} \text{ A}^2 \text{ m}^{-4} \text{ J}$, which can still be protected by an external resistor provided the magnet current is high enough. However, the magnet may be able to quench without any external quench-protection. To check this, use Equation 12 to estimate the decay time constant for the coil to be quench protected. For $J = 1.4 \times 10^8 \text{ A m}^{-2}$, $t = 12.8 \text{ s}$. As long as the coil quenches faster than this time constant, the magnet will be protected.

Let us now estimate the time for the entire one-meter section to become normal. From Equation 20, at a $J = 1.4 \times 10^8 \text{ A m}^{-2}$, the value of $V_L = 1.41 \text{ m s}^{-1}$. The time for the quench to propagate to the opposite side of the coil is 1.11 seconds. If the quench starts at the center of the magnet, the quench moves 0.5 meters at the rate of 0.0394 m s^{-1} . The time for the entire coil to go normal is 12.7 s. The time for the entire coil to go normal is roughly equal to the decay time constant that ensures that the hot spot temperature not exceed 300 K. Given the fact that the time for the entire coil to turn normal is equal to the desired time constant, it is likely that additional quench protection is needed. If the time for the entire coil to turn normal were 6 seconds, the coil would probably be self protected. A fully self-protected of a 1-meter long section of an infinite solenoid would be about 0.5 meters long, perhaps a little longer.

A low resistivity aluminum cylinder shrunk over the coil can protect the magnet by driving the entire magnet through quench-back. The aluminum cylinder must be closely coupled to the superconducting coil. (Closely coupled is a relative term. In this case the average radius of the aluminum cylinder should be less than 10 percent larger than the average radius of the coil for a coupling coefficient of 0.8 or larger.) The thickness of the ground plane insulation between the coil and the aluminum cylinder should be minimized, so that quench-back is not delayed.

One can calculate the effect of quench-back by calculating the minimum resistance R_{\min} needed to turn the superconducting normal is some time t_{QB} that is a fraction of the time constant for safe coil quenching. For many cases of interest for the muon cooling solenoids, t_{QB} should be about 0.25 times the time constant for safe quenching calculated using Equation 12. An expression for the minimum resistance needed to protect a quenching coil through quench-back is given as follows [12,16]:

$$R_{\min} = \frac{2\pi R N_1}{A_{c1} J_o} \frac{\rho_2 H_2}{t_{QR}}^{0.5} \quad -21-$$

where R is the radius of the coil, N_1 is the number of turns in the coil; A_{c1} is the cross-section area of the conductor; and J_o is the conductor current density at design current. H_2 is the enthalpy per unit volume of the shell material needed to increase the shell temperature from 4 K to 10 K; ($H_2 = 13200 \text{ J m}^{-3}$) ρ_2 is the resistivity of the shell material; and t_{QB} is previously defined as 0.25. A 1-centimeter thick RRR=25 aluminum shell will substantially shorten the time for a quench. As a result, the quench protection resistance can be reduced more than an order of magnitude. Quench-back allows coils to be self-protected.

Scaling of Large Solenoids Based on Quench Protection

From the standpoint of modeling stress in the coils and support members, building a scale model of a large superconducting solenoid magnet make sense particularly when the magnet is indirectly cooled with no liquid helium in the windings. From the previous section, we learned that in a scale model magnet the current density in the coils must go up as the scaling factor. In other words, a half-scale magnet must have coils with twice the current density of the full-scale magnet system. Quench protection is the other issue that must be looked at when doing magnet scaling. The following general rules should be applied when scaling a solenoid magnet:

The motivation for building a scale-model magnet is to be able to solve magnet construction, stress, and quench protection problems in a magnet system at lower cost. Building an operating scale-model magnet is useful when the solenoid magnet system is large and when the solenoid is indirectly cooled with either a cryocooler or helium flowing in tube attached to the coil structure. Scaling does not work very well for cryogenically stable magnet unless all that one is doing is measuring the stability criteria of the magnet.

The thickness of the stress carrying material in the coils and the support structure scale along with the solenoid coil radius. The radius that should be used is the average radius of the coil. The thickness and radii of all coils in a multiple coil system must scale together. The physical thickness of the coil may be different from what the scaling factor dictates because the insulation in the coils can not be scaled. The thickness of the cryostat vacuum shell components will also scale with coil radius.

The current density in the superconductor plus the surrounding matrix material is inversely proportional to the magnet scale factor. This means that the percentage of superconductor must go up in the scale model.

Scaling works best superconducting for solenoids that have a conductor with a large normal metal to superconductor ratio. If a full-sized magnet has ten percent superconductor ($r = 9$), the half scale model conductor will be 20 percent superconductor ($r = 4$) and a 40 percent scale model conductor will be 25 percent superconductor ($r = 3$). The superconductor limits the extent to which one can scale large magnets. It is recommended that the ratio of normal metal to superconductor for the reduced-scale model magnet not be below three (one-quarter superconductor). Reducing the normal metal to superconductor ratio to below three introduces stability problems in an impregnated indirectly cooled coil system. The normal metal to superconductor ratio in the coil material limits scaling more than any single factor.

Quench protection may be another limiting factor on scaling. A reduction of the normal metal to superconductor ratio in the conductor will increase the conductor hot spot temperature for a given current decay profile during a quench. In general, a reduced scale magnet will quench faster than the full-scale magnet. The time needed to get rid of the current in the coil during a quench is inversely proportional to the scaling factor squared. In general, the scale-model magnet is expected to perform worse in a quench situation than does the full-scale magnet.

Because there are quench implications introduced by scaling one has to ask a number of questions. Some of these questions include: 1) Should the magnet current scale as the inverse of the scaling factor or can it stay the same as the current in the full-scale magnet? 2) Should the scale model

solenoid scale only in radius (because the forces are primarily hoop forces) or should the model scale in all three dimensions? 3) Should the superconducting coil thickness scale with radius or should the thickness of the conductor scale with radius? The adjunct question is how far can one go in reducing the thickness of the insulation in the scale model coil without compromising the quench protection of the coil? 4) Does quench-back quench protection scale as quench protection using a dump resistor to extract the quench energy? These and other questions will be answered in Table 2 on the next page.

Table 2 compares three half-scale coils with a full-scale simple solenoid that is part of a long solenoid system. The full-scale solenoid has the following characteristics: 1) The radius to the center of the solenoid coil is 0.5 meters. 2) The length of the coil section is 1 meter. 3) The magnetic induction on axis is 5 T and the peak induction in the coil is also 5 T. 4) The conductor has a copper to superconductor ratio $r = 9$. (The conductor is 10 percent Nb-Ti.) 5) The coil thickness is 36.4 mm of which 32 mm is conductor and 4.4 mm is insulation between the 16 coil layers. 6) The design current for the full-scale solenoid is just over 1000 A, which is carried in a conductor cross-section of 8 square millimeters. 7) The J_c of the superconductor is assumed to be 2750 A mm at 4.2 K and 5T.

The three half-scale model coils shown in Table 2 have different characteristic and they demonstrate how the scaling process works. The three half-scale magnets have the following characteristics:

Half-Scale Model #1 has a coil radius to the center of the coil of 0.25 meters. The thickness of each conductor is reduced by a factor of two. The insulation thickness is the same as the full-scale coil, so the coil thickness is 20.4 mm instead of 36.4 mm (half of 36.4 mm). The length of the coil and the number of turns for this coil are the same as the full-scale coil.

Half-Scale Model #2 has the same radius and coil thickness as Half-Scale Model #1. The length of the coil is reduced by the scaling factor. As in the previous case, the thickness of the conductor is reduced a factor of two. The width of each conductor is also reduced by a factor of two. There are fewer turns in this coil because the coil insulation takes up more space in the length direction of this coil. The coil current is a little more than half of that of the full-scale coil and Half-Scale Model #1

Half-Scale Model #3 is the same as Half-Scale Model #2 except that the coil thickness is 18.2 mm instead of 20.4 mm. The coil thickness comes out of the superconductor matrix, so that the bare conductor thickness is 0.863 mm instead of 1.0 mm. This affects coil stress and the quench protection characteristics of the magnet.

Table 2 compare many parameters including the cost of the magnet. The cost factor for all of the magnets is based on studies done on a number of large detector solenoids. The cost equation is given as follows [17,18]:

$$C(M\$) = 0.53 [E(MJ)]^{0.70} \quad -22-$$

where C is the cost in millions of 2000 dollars and E is the stored energy in MJ. Other studies show a similar relationship between cost and stored energy. The cost equation above comes from studies that have been done on large detector magnets. The cost has been escalated to 2000 dollars. The cost model for large detector magnets given above includes the cost of the magnet cryostat and engineering. At the very least, the cost equation given above can be used to scale the material and fabrication cost for the solenoid magnet and its cryostat.

**Table 2. A comparison of Magnetic Stress and Quench Protection Behavior
For a Full-scale One-meter Diameter Solenoid and Three Half-scale Model Solenoids**

Parameter	Full Size	Half #1	Half #2	Half #3
Physical Parameters of a Magnet Section				
Induction on Axis (T)	5.0	5.0	5.0	5.0
Average Radius of Coil (mm)	500.0	250.0	250.0	250.0
Coil Section Length (mm)	1000.0	1000.0	500.0	500.0
Coil Physical Thickness (mm)	36.4	20.4	20.4	18.2
Number of Layers	16	16	16	16
Number of Turns per Layer	244	244	238	238
Number of Turns per Coil	3904	3904	3808	3808
Coil Packing Fraction*	0.858	0.765	0.747	0.722
Superconductor and Coil Parameters				
Bare Conductor Dimensions (mm)	2.0x4.0	1.0x4.0	1.0x2.0	0.863x2.0
Copper to Superconductor Ratio	9	4	4	4
Magnet Current (A)	1019.2	1019.2	522.4	522.4
Coil Current Density (A mm ⁻²)	109.3	195.0	195.0	218.6
Conductor Current Density (A mm ⁻²)	127.4	254.8	261.2	302.7
S/C Current Density (A mm ⁻²)	1274	1274	1306	1514
Magnet Section Stored Energy** (MJ)	7.813	1.953	0.977	0.977
Magnet Section Self Inductance (H)	15.04	3.761	7.157	7.157
Magnet Stress and Strain Parameters				
Ave. Coil Stress w/o Shell (MPa)	159.3	159.3	163.2	180.1
Ave. Coil Strain w/o Shell (%)	0.107	0.107	0.109	0.121
Ave. Coil Stress with Shell [^] (MPa)	134.0	134.0	137.3	147.8
Ave. Coil Strain with Shell [^] (%)	0.090	0.090	0.092	0.099
Peak Field BJR Stress (MPa)	306.9	305.5	313.2	364.6
Quench Protection Parameters				
Safe Quench Time Constant (sec)	16.08	3.57	3.40	2.53
Dump Resistance No QB (ohm)	0.935	1.054	2.105	2.829
Dump Voltage No QB (V)	953	1074	1100	1478
EJ ² Limit (A ² m ⁻⁴ J)	1.27x10 ²³	1.27x10 ²³	0.67x10 ²³	0.90x10 ²³
Quench Propagation Velocity (m s ⁻¹)	1.20	3.90	4.07	5.23
Time for Section to go Normal (sec)	>10.7	>2.8	>2.1	>1.6
Self-Protecting w/o Quench Back	Maybe	Maybe	Maybe	Maybe
R for QB from RRR=25 Al [^] (ohm)	0.030	0.058	0.112	0.110
Dump Voltage for Quench-Back [^] (V)	30.5	59.1	58.5	57.5
Self-Protecting with Quench-Back [^]	Yes	Yes	Yes	Yes
Coil Section Cost Factor ^{^^} (Full Size = 1)	1.00	0.38	0.23	0.23

* The turn to turn insulation is 0.1 mm thick in all cases. The layer to layer insulation is 0.275 mm in all cases.

** This is the self-inductance of the section, which is calculated from the stored energy for that section.

[^] An aluminum cylinder (RRR=25) outside of the coil is used for quench back and to carry magnetic forces. The full-scale cylinder is 12.7 mm thick; the half-scale cylinder is 6.35 mm thick.

^{^^} The cost of the solenoid section is proportional to the stored energy to the 0.7 power.

Table 2 on the previous page compares the stress and quench protection parameters for three half-scale solenoids compared to a full-scale 5 T solenoid with a diameter of 1 meter (at the center of the coil) and a length of 1 meter. In both the full-scale and half-scale cases, the solenoid is assumed to be part of a very long string of similar solenoids. For the case of the half-scale #1, the overall coil package current density is 1.78 times the current density of the full-scale solenoid, yet the current density in the conductor of the half-scale model is twice that of the full scale model. The difference is the fact that the insulation thickness is the same for the half-scale and full-scale magnets. In half-scale #2, the conductor current density is 2.5 percent larger than it is in half-scale #1 case. This difference is due to the fact that the bare conductor length was reduced by a factor of two. The conductor current density for half-scale #3 is about sixteen percent larger than for half-scale #2. This difference is due to the fact that conductor had to be removed so that the thickness of the coil in half-scale #3 is exactly half of the coil thickness for the full-scale case.

Cases half-scale #1 and #2 have similar current densities. In both cases, the average stress and the peak BJR stress are within a 2.5 percent of these same stresses in the full-scale case. Case half-scale #3 is different, because the stresses are higher by nearly twenty percent. It is clear that from a stress standpoint, that either case half-scale #1 or half-scale #2 are superior to case half-scale #3. The conclusion is that one should reduce the thickness of the conductor in the coil not the thickness of the coil itself, when one does the scaling. This is true, even if one puts a metal cylinder on the outside of the coil to reduce the average stress in the winding and induce quench-back in the winding.

The aluminum cylinder thickness was 12.7 mm in the full-scale and 6.35 mm for the three half-scale cases. The effect of the aluminum cylinder was to reduce the coil stress and strain by about 19 percent. If the aluminum had the same elastic modulus as the copper based conductor, the stress reduction in the coil would be about 40 percent. If the cylinder on the outside of the coil is used to reduce the stress in the coil alone, that cylinder should be made from a high modulus material. The 1100-O aluminum (RRR=25) cylinder also induces quench back into the coil. A stainless steel cylinder can not do that. When modeling the magnet using scale models, all support-cylinders and other stress carrying structures must be made from the same material as the full-scale magnet.

The time constant for safe quenching for the full-scale coil is about 16.1 seconds. Perfect scaling would require that the scale model coils have a time constant of about 4 seconds. The reduced copper to superconductor ratio and the increased current density in the conductor will reduce the safe quenching time constant. Cases half-scale #1 and half-scale #2 show a reduction of more than four but the values given are not too far off. When one looks at the dump voltage for safe quenching the voltages calculated for cases half-scale #1 and half-scale #2 are within fifteen percent of the full-scale case. It is clear that the half-scale cases that are of interest are slightly worse from a quench standpoint than the full-scale case. Scaling as was done with the half-scale #3 case is not recommended.

The resistance needed to induce quench back was calculated using equation 21 then they were multiplied by a factor of three to account for imperfect coupling and the time constant needed for the heat to cross the ground plane insulation between the cylinder and the coil. The quench back resistance varies with the radius of the coil. This is a case where the half-scale magnets do not quench-back as easily as the full-scale magnet. It is a fact of life that large diameter solenoid quench-back easier than small diameter solenoids. This is why quench-back has been used to protect large detector magnets.

In general, if the smaller scale magnets are designed to quench safely, the full-scale magnet will also quench safely. From the standpoint of quenching, building a scale model magnet does not cause problems when extended to the full-scale magnet. From the standpoint of quenching and stress using scale models is a cost-effective way to develop superconducting solenoids that can be used as part of the muon collider or neutrino factory muon-cooling system.

The field Generated by Short Coils and the Gaps in Long Coils

The muon-cooling channel consists of a series of short coils that generate field like a long coil. It is useful to look at the field generated by short coils alone. A series of short coils can be combined to simulate long coils. A gap in a long coil can be simulated by adding the field from a long coil to the field generated by a short coil operating at opposite polarity.

The magnetic induction at a position z on axis for a short solenoid of radius R and a current per unit length of nI can be given by the following analytic expression provided the solenoid is thin compared to its radius [4]:

$$B_z = \frac{\mu_0 n I}{2} [\cos \beta_2 - \cos \beta_1] \quad -23-$$

where B_z is the magnetic induction on axis; μ_0 is the permeability of air ($\mu_0 = 4 \times 10^{-7} \text{ H m}^{-1}$); and nI is the ampere per unit length for the solenoid. The angles β_1 and β_2 are defined as follows:

$$\beta_1 = \tan^{-1} \frac{R}{z - z_1} \quad -23a-$$

and

$$\beta_2 = \tan^{-1} \frac{R}{z - z_2} \quad -23b-$$

where R is the radius of the thin solenoid; z_2 is the z position of one end of the solenoid; z_1 is the z position of the other end of the solenoid; and z is the point on the axis where the magnetic induction is calculated. On the axis of the solenoid, there is no r component of the induction, so B_r is zero.

Inside of the solenoid either β_1 or β_2 is negative. The maximum value of the induction B_z generated by a short coil will occur at $z = (z_1 + z_2)/2$, at the center of the coil. The absolute value of the induction at the center of a short solenoid B_{zc} is:

$$B_{zc} = \mu_0 n I (\cos \beta_o) \quad -24-$$

where

$$\beta_o = \tan^{-1} \frac{2R}{L} \quad -24a-$$

where R is the radius of the coil and L is the length of the coil.

The field for short coils was generated as a function of z for two coils. The large coil has a value of $R = 0.5$ meters and a value of $L = 0.2$ meters. The small coil has a value of $R = 0.25$ meters and a value of $L = 0.2$ meters. Both coils have a value of $nI = 3.9789 \times 10^6 \text{ A m}^{-1}$, which in a long solenoid would generate an induction of 5 T. If the coil were 25 mm thick, the current density in the coil would be 159.2 A mm^{-2} . Figure 4 on the next page shows the value of the magnetic induction for the large coil and small coil as a function of z . The center of the short coils is at $z = 0$.




Figure 5. The Magnetic Induction on Axis versus z Position for Two Short Coils

Figure 5 above shows the field generated by two 0.2-meter long solenoidal coils. These coils are thin and carry a current of 795780 A. The large coil has a current sheet radius of 0.5 m; the small coil has a current sheet radius of 0.25 m. From Figure 5, the peak induction on axis is larger for the small coil than it is for the large coil. However, the zone of influence for the large coil occurs over a longer distance than it does for the small coil.

The expression for the self-inductance L for a thin solenoid of length Z , radius R , and number of turns per unit length n is given by the following analytical expression [4,19]:

$$L_1 = \pi \mu_0 R^2 n^2 [(Z^2 + R^2)^{0.5} - R] \quad -25-$$

For a given solenoid coil current I , the stored magnetic energy can be calculated using the following expression:

$$E = \frac{L_1}{2} I^2 \quad -26-$$

For long coils where $Z \gg R$, the self-inductance of the solenoid is proportional to R^2 and Z when n is constant. For short solenoids where $R \ll Z$, the self-inductance is proportional to R and Z^2 , when n is constant. For the large solenoid in Figure 5 with $n = 5000$, $L_1 = 0.9504$ H. For the small solenoid in Figure 5 with $n = 5000$, $L_1 = 0.4325$ H. If both the large and the small coils have a current of 795.8 A (enough to generate 5.0 T in a very long solenoid of either radius with 5000 turns per meter), the stored energy of the large and small coils would be 475 kJ and 216 kJ respectively.

The short coil studies shown in Figure 5 suggest that if one wants to influence the field on axis over a relatively short length, one should use a small diameter coil to influence the field. Fields generated by short small diameter coils will be higher and the stored energy generated by that coil is smaller. Since, the stored energy of these coils is smaller, their fabrication cost will be lower as well.

The converse of the problem of the field generated by a short coil is the field (or lack of field) generated by a gap in a long coil. Figure 6 below is a plot of the magnetic induction on either side of a long solenoid that is 8.2 meters long and 1 meter in diameter. The solenoid current per unit length is the same as would be needed to generate a uniform induction of 5 T in an infinitely long solenoid. Over a length from -1.5 m to $+1.5$ m about the center of the solenoid, the induction droops as one moves out from the solenoid center. (There would be no droop in the gap = 0.0 case if the solenoid was infinitely long.) The effect of gaps of 0.2-meters, 0.4-meters, and 0.6-meters is shown in Figure 6. When one compares Figure 5 and Figure 6 one can see that the drop in the field due to a 0.2-meter gap is equal to the induction generated by the large coil, which is 1.0 meter in diameter and 0.2 meters long.




Figure 6. The Effect of Various Gaps in the Coil of a Long Solenoid on the On Axis Field

A 0.2-meter long gap in a 1-meter diameter long solenoid will reduce the field on axis about twenty percent. The 0.4-meter long gap reduces the field on axis about thirty-eight percent. It is clear that the gap between coils should be minimized in terms of perturbation of the magnetic field on axis, because there could be of some consequences in a muon-cooling channel. One of the reasons for having gaps between coils is the RF feed to the 200 MHz RF cavities. The size of the gap between the coils is dictated by the RF wave-guide dimensions and the minimum space needed for the cryostat and thermal insulation. For 200 MHz RF cavities, the minimum gap between the coils is between 0.2 and 0.3 meters.

An alternative approach is to build-up the coils at the ends on either side of the gap. The built-up length at the end would be half the gap width. If the gap is 0.4 meters, the built-up portion of the coils would double the coil thickness (or current density) for 0.2 meters on either side of the gap. Figure 7 below compares the effect of a 0.4 meter gap on the on axis induction of a 1.0 meter diameter solenoid with and without a build-up at the ends of the coils.




Figure 7. The Effect of Building-Up the Ends Around a 0.4-meter Gap in a Solenoid

Figure 7 above shows that building up the ends on either side of the gap between two solenoids improves the quality of the magnetic field on axis. A 0.4-meter gap in a 1.0-meter diameter solenoid caused a drop in the field of 38 percent. Building up the coils ends on both sides of the gap reduces the field drop to about 12 percent. The integrated RMS field error along the axis is even lower. The downside of built-up coil ends around a gap in a superconducting solenoid is that the peak magnetic induction in the conductor is increased.

Large and Small Solenoids with Periodic Gaps

The muon cooling system consists of a periodic structure of 200 MHz RF cavities and liquid hydrogen absorbers. Both the RF cavities and the absorbers require radial access through the superconducting solenoid coils. The size of the gap needed to provide these services varies from 0.1 to 0.15 meters for the hydrogen absorber cryogenic services to 0.2 to 0.25 meters for the RF cavity wave-guides. One approach is to make a periodic magnet structure where the coils and gaps are periodic and the gap length is equal to the coil length.

The effect of having a periodic field in the solenoid channel may be to excite a resonance or two in the beam causing a significant number of particles to be lost before they reach the end of the cooling channel. The number of particles lost appears to be a function of the following parameters: 1) the momentum of the particles that one wants to pass through the magnet structure, 2) the period of the field variation, and 3) the amplitude of the field variation. For cooling, the momentum range for the muons is from 150 MeV/c to about 225 MeV/c depending on where one is in the cooling channel and the quality of the phase-rotation before cooling starts. Even if the particles have a very low momentum spread, there will be a minimum momentum variation of something like ± 10 MeV/c for the muons.

The study of particle resonance phenomena is not the purpose of this report. Ways that reduce the variation of the magnetic induction on axis are very much a part of this report. To show the effect of gap and coil radius on field uniformity, two cases were studied. The first is large solenoid (radius 0.5-m) with coil lengths and gaps between coils that are 0.2-meters long. The second is a small solenoid (radius 0.25-m) with coil lengths and gaps between coils that are 0.2-meters long. Both cases have coils that carry the same number of ampere-turns per meter as would be needed to produce a 5 T in a very long solenoid. (For an induction of 5.0 T in a long solenoid, 3.9789 MA m^{-1} are needed.) In both cases, the average induction on axis is 2.5 T




Figure 8. A comparison of Small Diameter and Large diameter Solenoids With Alternating 0.2 Meter Long Coils and Gaps

Figure 8 compares the on axis field variation for the two cases where a solenoid system has coils that are 0.2 meters long and the gap between the coils is 0.2 meters. The large solenoid has an average diameter of 1.0 meters; the small solenoid diameter is 0.5 meters. The current in the coils per unit length is the same as the current per unit length that generates 5.0 T in an infinite solenoid. Since half of the space along the length is gap, the average induction on the solenoid axis should be 2.5 T. The average induction on axis is not quite 2.5 T, because the solenoid length is finite (about 8 meters). When the coil diameter is 1.0 meters, the effect of the gaps between coils is a field variation of ± 0.16 percent. When the coil diameter is cut in half to 0.5 meters, the field variation due to the same gaps is increased to ± 6.78 percent.

A solenoid with alternating coils and gaps produces a more uniform field than a solenoid that is continuous with a single gap of the same dimension. From a field uniformity standpoint, it appears that a solenoid with a series of alternating coils and gaps is a desirable thing for the cooling channel. The gaps will allow access to the RF cavities and the hydrogen absorbers. If one want to produce an induction of 5 T on axis, the current density or the coil thickness must be doubled. The down side of the alternating coils and gaps is the fact that the peak field in the superconducting winding must go up compared to the continuous solenoid case. The field rise in the coil depends on the coil thickness compared to the radius and a number of other factors. The factor that should be considered is the effect of the periodic field variation on the stability of the beam. In other words, do field periodicities produce a resonance that kicks out muons in the desired momentum range?

Large and Small Solenoids that Flip the Field Polarity

A feature common to most of the cooling channels studied is flipping the field in the channel from one polarity to the other. The FOFO and Super FOFO channels have field that changes polarities every few meters (See Figure 1). The single flip channel changes polarity only once or twice as the muons go down the cooling channel. The flip sections require that the coils generate the magnetic gradient. The larger this gradient is, the larger the forces between coils are for a given coil radius. In order to minimize magnet cost, it is important to correctly place the coils that generate the field gradient in the flip zone.

Two cases of flip configuration were studied with both large and small diameter coils. The First configuration is similar to the FOFO configuration where two 0.2-meter long coils are charged to one polarity and then two more coils are charged to the opposite polarity and so on. The gap between the coils at the same polarity is 0.2 meters and the gap between coils in the flip region is also 0.2 m. The second configuration has 0.2-meter long coils of alternating polarity. The gap between the coils in this case is also 0.2 meters. A schematic representation of the two cases is shown in Figure 9. In both cases studied, the average diameter R of the large coil is 1.0 meter and the small coil average diameter R is 0.5 meters. The coil lengths L and the gaps between coils G are the same for both cases regardless of the value of the coil radius R .

The graph of the induction on axis for the first case, the two coils of one polarity and two coils of the opposite polarity, is shown in Figure 10. The graph of the induction on axis for the second case, the alternating polarity case, is shown in Figure 11.

The thickness and current density of the coils shown in Figures 9, 10 and 11 are the same as one use would to generate an induction of 5 T in a long continuous solenoid. (For an induction of 5.0 T in a long solenoid, 3.9789 MA m^{-1} is needed.) Because there is a gap between coils G that is equal to the length of the coils L , the highest average on-axis induction one can expect would be 2.5 T as is shown in the case given in Figure 8. One sees a very different result, when one compares Figures 10 and 11 with Figure 8.





Figure 9. A schematic Representation of Two Flipped Polarity Solenoid Cases




Figure 10. The On-axis Induction versus Distance along the Axis For the Case where Polarity Changes Every Two Coils




Figure 11. The On-axis Induction versus Distance along the Axis for the Case where Polarity Changes Every Coil


The case where the polarity changes every two coils yields some interesting results (shown in Figure 10). The period of the field variation of both the large coil case and the small coil case is 1.6 m (the periodicity of the alternating polarity). The on axis field variation for the large coil case (with an average coil diameter of 1-meter) looks sinusoidal. The on axis field variation for the small coil case (with an average coil diameter of 0.5 m) does not look sinusoidal. The small coil case generates an on axis gradient of over 11 T m^{-1} and a peak on axis induction of about 1.8 T. The large coil case has an on axis gradient of less than 5 T m^{-1} and a peak on axis induction of just under 1.0 T. There is a structure of the field generated by the small coils in the region where both coils have the same polarity. This structure is similar to the structure of the small coil data shown in Figure 8. From Figure 10, it is clear that the current density or the coil thickness must be increased in order to achieve a peak induction on axis of 2.5 T. For the large coil case the coil thickness must be increased by a factor of 2.5. For the large coil case, the induction in the conductor is above 7 T when the peak induction on axis is 2.5 T. The small coil thickness need only be increased by a factor of 1.39 to achieve an on axis peak induction of 2.5 T. In this case, the induction in the coil conductor approaches 5.T. The cost for the large coil case is dictated by the diameter of the coils and their thickness. To first order, the cold mass for a solenoid per unit length is proportional to the coil average radius and its thickness. The solenoid cost is proportional to its mass to about the 0.7 power. For a given solenoid design the coil thickness goes up with radius. For a given solenoid radius, changes in design will affect the coil thickness hence the magnet cost.

The case where the polarity flips with every coil also yields interesting results (see Figure 11). In this case the on axis field period is 0.8 m. Both the large coil and the small coil cases look sinusoidal. The peak on axis gradient for the small coil case is about 9 T m^{-1} , whereas the peak on axis gradient generated by the large coils is about 2 T m^{-1} . The peak on axis induction for the small coil case is 1.4 T compared to a peak on axis induction of about 0.35 T for the large coil case. For the small coil case, the thickness of the coil must be increased a factor of 1.79 in order to achieve a peak on axis induction of 2.5 T. For the small coil case, with a peak induction of 2.5 T on axis, the induction in the


superconductor would be over 6 T. As a comparison, the large coil thickness must be increased by a factor of about 7.2 in order to achieve the same on axis induction of 2.5 T. The peak induction in the coil superconductor would be well above 15 T in order for the induction on axis to increase to 2.5 T. Clearly, from a cost standpoint rapid flipping of the field with large diameter coils is not acceptable.

The Geometry of Flipping Field Coils for Minimum Cost

The solenoids that change the polarity of the magnetic field on axis cost money if they are properly designed to minimize stress and the peak induction in the superconductor. This section shows some designs that move toward minimizing the cost of solenoids that flip the magnetic field. As an example, a single flip section was studied using the 1-meter diameter interrupted solenoid (a 0.2-meter long coil and gap) presented earlier. Studies were made for a single flip section with large coils and a combination of large and small coils is presented in this section.



a) *Single Field Flip with Large Coils Only*



b) *Single Flip with Large Plus Small Coils*

Figure 12. A Schematic Representation of Two Single Flip Coil Systems

Figure 12 above shows two coil configurations for flipping the field polarity at a single point in a string of solenoid coils. The coils leading into the flip region are 0.2-meter long coil alternating with 0.2-meter long gaps. The current in the coils is 3.9789 MA per meter, which in a long solenoid would develop an induction of 5 T on axis. With the alternating coils and gaps, the on axis induction is 2.5 T

(see Figure 8). The average radius of the coils coming into and out of the field flip region is 0.5 m. In the field flip region, the coils can be smaller as is shown in the b part of Figure 12.

The spacing S between the coils where the polarity flips determines the peak induction in the superconductor and the magnetic forces that are pushing the coils apart. A value of S that results in little or no field rise in the superconductor can be calculated the following expression:

$$A_f = 2\pi R^2 = 2\pi RS \quad -27-$$

where A_f is the flux carrying area; R is the average radius of the coils; and S is the space between the coils of opposite polarity. Solving Equations 27 for S , one finds that $S = R$ when there is no increased induction at the gap between the two coils. In reality there is an increase in the field in the coil next to the flip region due to the graininess of the coil structure. Therefore it is not unreasonable to reduce the value of S to 0.8 or 0.9 R . We shall see from Figure 13 on the next page, reducing S below 0.8 R has only a small effect on the gradient generated by the flipped polarity in the coils. Values of S substantially below result in a much higher induction at the conductor and larger forces that are trying to push the coils of opposite polarity apart. Both effects will increase the cost of the solenoid in the reversed polarity region.




Figure 13. The Effect of Coil Spacing S in the Flipped Polarity Region in a 1 meter Diameter Solenoid with the Flip Coil $R = 1.0$ m

Figure 13 on the previous page show the effect of the spacing S of the reversed polarity coils on the field generated by the solenoid in the field reversal region. A long distance from the field reversal point at $x = 0$, the induction is either $+2.5$ T or -2.5 T. The optimum gap from a geometric perspective is $S = R = 0.5$ meters. At a value of $S = 0.2$ m, the induction gradient at $x = 0$ is 5.26 T m^{-1} . When $S = 0.4$ m, the gradient at $x = 0$ is 4.85 T m^{-1} . $S = 0.6$ m, produces a gradient at $x = 0$ of 4.01 T m^{-1} . When $S = 0.8$ m, the gradient at $x = 0$ is 3.10 T m^{-1} , and when $S = 1.0$ m, the gradient at $x = 0$ is 2.31 T m^{-1} .

One can increase the gradient at the flip point by increasing the current density in the coil or by increasing the coil thickness while keeping the current density constant. One can also increase the gradient at the flip point by using smaller diameter coils to create that gradient. The smaller diameter coils can be closer together without generating large forces between the coils and raising the peak field in the conductor.




Figure 14. A comparison of Induction on Axis Versus x for Flux Reversal Sections Using Large coils $R = 0.50$ m and Small Coils $R = 0.25$ m

Figure 14 above compares the flux reversal section where flux reversal is done using the large solenoid coils $R = 0.50$ meters and a pair of small coils $R = 0.25$ meters to reverse the flux in the flux reversal region. In both cases, the solenoids entering and leaving the flux reversal region have an average diameter of 1.0 meter and the solenoid consist of alternating 0.2 -meter long coils and 0.2 -meter long gaps. The current in the coils is 3.9789 MA per meter, which would generate a uniform induction of 5.0 T in a long continuous solenoid. The induction well away from the flux reversal section in both cases is 2.5 T. The length of the flux reversal coil on either side of $x = 0$ is 0.2 meters. These coils are separated from the standard solenoid by a gap of 0.2 meters in the x direction. The flux reversal coil gap S for both cases is $0.8 R$.

Reducing the radius of the flux reversal coil greatly increases the gradient in the flux reversal region. The gradient for the $R = 0.5$ m case at $x = 0$ is 4.85 T m^{-1} . Reducing the flux reversal coil radius R to 0.25 m, increases the gradient at $x = 0$ to 12.35 T m^{-1} . A factor of 2.55 increase in the $x = 0$ gradient was generated by reducing R by a factor of 2.

The potential down side of simply reducing R is the fact that the field transition in the flip region is no longer smooth. (From the standpoint of a muon-cooling channel, the dip in the field on both sides of the flux reversal region may not be bad.) One can control the gradient and the field shape on either side of the flux reversal region by changing the gap S between the flux reversal coils, or by changing the current in the flux reversal coils. One can also change the field shape by varying the radius of the coils coming into the flux reversal region and by varying the gap between the flux reversal coils and the main solenoid coils. Figure 15 below shows the effect of changing the current in the $R = 0.25$ -meter coils in the flux reversal region.




Figure 15. The Effect of Changing the Current in $R = 0.25$ m Flux Reversal Coils On the Magnetic Induction Profile as a Function of Distance from the Flux Reversal Point

From Figure 15, one can see that changing the current in the coils in the flip region will change the gradient of the field and the shape of the field on either side of $x = 0$. When the small coils carry $3.98 \text{ MA per meter}$ the gradient is 12.35 T m^{-1} . When the current in the coils is reduced to $2.65 \text{ MA per meter}$, $1.33 \text{ MA per meter}$ and zero, the gradients at $x = 0$ in the flip region are 9.00 T m^{-1} , 5.66 T m^{-1} , and 2.31 T m^{-1} respectively. The zero current case is the same as the large coil case where $S = 1.0$ -meter. It is clear from Figure 15 that more than current manipulation in the small coils is needed to smooth the field and maintain the gradient the smaller flux reversal coils allow. Other approaches for smoothing the field include varying the diameters of the coils coming into the region and changing the spacing between coils.

Concluding Comments on Muon Cooling Solenoids and How to reduce their Cost

All of the design studies done for this report were done for 1.0-meter and 0.5-meter diameter coils. Most of the muon cooling work for the neutrino factory cooling system requires that the coils have an average diameter of about 1.5-meters. The size of the largest solenoid for the neutrino factory cooling system is dictated by the RF cavities that are inside of the coils. The selection of the RF frequency was based on the size of the RF bucket needed to accept the muons coming from the phase rotation system. Most of the neutrino-factory cooling studies have assumed 200 MHz RF cavities. These cavities require solenoids with a warm bore that is about 1.2 meters in diameter. The minimum inner radius for the superconducting coils is about 0.7 meters; hence the average coil radius is close to 0.75 meters.

For a given induction on axis, the cost per unit length of a superconducting solenoid around the RF cavities is proportional to the average radius to the 1.4 power. For a given radius, the cost per unit length is proportional to the magnetic induction to the 1.4 power until one must change superconductor technology. In order to reduce the magnet cost per unit length, one must reduce the magnet coil average radius and minimize the increased coil thickness that is caused by rapid changes in the magnetic field on axis. In order to minimize the cost of the solenoid magnets, one must move in the direction of reducing the mass of the superconducting coils and the surrounding support structure, which mean that BJR stress limits will be pushed within the coils.

To be more specific, the following design steps should be utilized to minimize the cost of the superconducting solenoid that are an integral part of the muon-cooling channel:

The RF frequency should be maximized in all parts of the cooling channel. Low RF frequencies mean large RF cavities, which increases the cost of both the cavities and the solenoid that surrounds them. The RF frequency chosen is based on the acceptance of muons from the phase rotation system. If acceptance can be improved, the RF frequency can be increased. If possible, the RF frequency should be increased as one moves down the cooling channel. This permits the magnets for later cooling stages to be smaller and it permits more efficient cooling because the solenoidal field can be increased for a given amount of money spent.

If possible, incorporate the solenoid coils inside low frequency RF cavities. Whether this is possible depends on whether beam is lost due to a larger variation in the magnetic field along the axis. Resonance studies would be useful for determining the momentum of the particles that are lost from the channel.

Reduce the induction on axis where the solenoid coils are the largest. This principle was employed in the cooling channel for the muon collider, and it should be employed in the neutrino-factory-cooling channel as well. Minimizing the cooling channel stored magnetic energy will result in a less expensive solenoid system for that channel. This is true even if the cooling channel is longer.

Short coils with equal length gaps produce a more uniform field than do long coils with the same short gap. Short periodicity of the magnetic field on axis may be useful from a beam resonance standpoint. Since the RF cavities require that there be gaps between coils, design the solenoid coils with periodic gaps. For a given uniform induction on axis, the periodic coil solution costs only a little more than the continuous solenoid solution provided the solenoid superconducting coils are relatively thin compared to their radius.

Do not make rapid changes of the magnetic induction on axis using large coils. This statement is particularly true when the channel has continuous changes in magnet polarity. When possible, use small coils to do polarity flips. This means that one should not change the polarity in channel magnets around large RF cavities. Smaller coil will produce larger gradients on axis without increasing the field at the conductor. From a magnet cost standpoint, this means that flux reversals in the cooling channel should occur at or near the hydrogen absorbers.

Do not use a gap between coils that change the polarity of the field that is much smaller than the coil average radius. Coils closer together produce only slightly more gradient on axis than coils that are an average radius apart. The magnetic induction at the conductor goes up and the forces pushing the coils apart go up without any real gain in magnet performance on axis. Using gaps that are smaller than 0.8 times the average coil radius is like throwing money away.

Use high current density coils to produce the magnetic field on the axis. The current density in the windings is dictated by winding stress and quench protection. The BRJ stress is less important if high local coil stresses can be carried by stiff adjacent conductors or a stiff support structure. Cryogenically stable coils are not very stiff in the radial direction. As a result, the BJR stresses become more important. In general, high current density coils will cost less to fabricate than lower current density coils. For the muon collider cooling channel, cable in conduit coils are probably not cost effective for a typical cooling channel solenoid.

The use of niobium-tin conductors should be avoided. It is better to cool Nb-Ti using helium at 2 K than it is to use niobium tin conductor. When niobium tin is used, it should be used in the region of the coils where its use is clearly called for. If niobium tin insert solenoids are used inside of niobium titanium solenoids the high field point for the magnet should be designed to be near the center of the inner bore of the niobium tin magnet.

Use a low mass-density conductor and a low mass-density support structure for large solenoids that operate at moderate fields on axis (say less than 4 T generated when using Nb-Ti as the superconductor of choice). Aluminum matrix conductors combined with aluminum support cylinder may be the cost-effective choice for many of the magnets in the neutrino-factory cooling channel. The coils made from aluminum matrix conductor should be inside of the aluminum support shell. An aluminum support shell can be used with coils fabricated from copper based superconductors, but aluminum based superconducting coils can not be used inside of a stainless steel support shell. Reducing the mass of the solenoids in the cooling channel will result in a reduction in the overall cost for the magnet system.

Avoid the use of bath-cooled and cryogenically stable magnet structures. Helium bath cooled coils are more compliant in the radial direction, hence magnetic forces applied to the inner windings of the solenoid are not transmitted to the outer windings. The helium bath cryostat adds to the mass of the magnet system and to its cost. Helium bath cryostat vessels must meet the ASME pressure vessel code, which puts additional restrictions on the magnets and their placement in the cooling channel. The magnet coils should be cooled by conduction from the support shell, if possible. The support shell can be cooled using two-phase helium in cooling tubes attached to the shell [20].

The preceding set of recommendations point toward using potted superconducting coils for the muon cooling channels. Potted coils in some parts of the muon collider may be subject to radiation damage. This should be investigated. Potted solenoid coils will some require development. This means that solenoid modules will have to be built before one builds the cooling channel solenoids using this technology. Since the muon-cooling channel is a series of repeating elements, this development is warranted in order to reduce the overall cost of the channel.

Since the solenoids in the cooling channel are relatively large, a scale model of the typical solenoid cell can be fabricated as part of the development process. If the scale model magnet operates with no training below its design current, the full-scale magnet should operate without training as well. A scale model can be used to solve stress and quench protection problems that can occur on the full-scale magnet. For early development of the solenoid system for a muon-cooling channel, fabrication of scale model magnets can be very cost effective. Scale models work well if the matrix to superconductor ratio in the full-scale magnet system is large. Scale model magnets may not be cost effective for modeling niobium tin magnets that will tend to have a smaller normal metal to superconductor ratio in the conductor.

Acknowledgements

The author acknowledges discussions he has had with Harold Kirk of the Brookhaven National Laboratory concerning the resonances generated by varying magnetic field along the axis of a phase-rotation channel. The author acknowledges discussion that he has had with Mike Zisman of the Lawrence Berkeley National Laboratory concerning whether it makes economic sense to build smaller-scale solenoids rather than full-scale solenoids when developing the solenoids for a muon cooling channel. The author acknowledges discussion with Dan Kaplan of IIT, Jocelyn Monroe of Fermilab and Robert Palmer of the Brookhaven National Laboratory concerning various types of solenoids for a muon-cooling channel. The author acknowledges discussions with J. R. Miller, Y. Eyssa, S. Kenny and S. Prestemon of the National High Magnetic Field Laboratory concerning the stress and strain issues in large superconducting solenoids. Finally the author acknowledges discussion he has had with Bert Wang concerning the design of large potted solenoid that can be indirectly cooled using either a cryocooler or two phase helium flowing in tubes attached to the coil packages.

This work was performed at the Lawrence Berkeley National Laboratory with the support of the Office of Science, United States Department of Energy under the DOE contract with the University of California number DE-AC-03-76SF00098.

References

The following list of references may be useful for those who want to look seriously at the design of a large superconducting solenoids that may be part of a cooling channel for a muon-collider or neutrino-factory.

- [1] Talks given at the Muon Collider Collaboration Meeting, 17-19 May 2000, Catalina Island CA by Dan Kaplan and Jocelyn Monroe
- [2] A talk given by Robert Palmer at the Muon Collider Collaboration Meeting, 17-19 May 2000, Catalina Island CA concerning the super FOFO cooling lattice.
- [3] J. R. Miller, National High Magnetic Field Laboratory, Tallahassee FL, private communication concerning the build up of stress at the high field point of thick superconducting solenoids
- [4] W. R. Smythe, Static and Dynamic Electricity, McGraw-Hill Book Company Inc., New York (1950)

- [5] E. P. Popov, Mechanics of Materials, Prentice-Hall Inc., Englewood Cliffs NJ, (1952)
- [6] R. J. Roark and W. C. Young, Formulas for Stress and Strain, 5th Edition, McGraw Hill Book Company, New York, (1975)
- [7] M. A. Green, "High Energy Physics Particle Detector Magnets," Wiley Encyclopedia of Electrical and Electronics Engineering, Volume 8, p739-752, A Wiley-Inter-science Publication, John Wiley and Sons, New York (1999)
- [8] M. A. Green, "Calculating the J_c , B, T Surface for Niobium Titanium Using the Reduced State Model," IEEE Trans. MAG-25, No. 2, p 2119, (1989)
- [9] B. J. Maddock and G. B. James, Protection and Stabilization of Large Superconducting Coils, Proceedings of IEE **115**, No. 4, p 543 (1968)
- [10] P. H. Eberhard, et al, "Quenches in Large Superconducting Magnets," Proceedings of the 6th International Conference on Magnet Technology, Bratislava Czechoslovakia, p 654 (1977)
- [11] M. A. Green, "Quench Protection and Design of Large High-Current Density Superconducting Magnets", IEEE Transactions on Magnetism **MAG-17**, No. 5, p. 1793, (1981)
- [12] M. A. Green, "Design Conditions for Fail-safe Quenching of a High-Current Density Superconducting Solenoid Magnet with a Shorted Secondary Circuit", LBL-14859, (1982)
- [13] P.H. Eberhard et al., "The Measurement and Theoretical Calculation of Quench Velocities Within Large Fully Epoxy-Impregnated Superconducting Coils", IEEE Transactions on Magnetism **MAG-17**, No. 5, p. 1803, (1981)
- [14] M. A. Green, "PEP Detector Development, Large Superconducting Solenoid, Copper Stabilized vs. Aluminum Stabilized Superconductor," LBL Engineering Note M5044A, January 1978
- [15] M. Scherer and P. Turowski, "Investigation of the Propagation Velocity of a Normal Conducting Zone in Technical Superconductors," Cryogenics **18**, p. 515 (September 1978)
- [16] M. A. Green, "Quench Back in Thin Superconducting Solenoid Magnets," Cryogenics **24**, p 659 (1984)
- [17] M. A. Green, R. A. Byrns, and S. J. St. Lorant, "Estimating the Cost of Superconducting Magnets and the Refrigerators Needed to Keep Them Cold," Advances In Cryogenic Engineering **37**, p 637, Plenum Press, New York, (1991)
- [18] M. A. Green and S. J. St. Lorant, "Estimating the Cost of Large Superconducting Thin Solenoid Magnets," Advances in Cryogenic Engineering **39**, p 271-276, Plenum Publishing, New York, (1993)
- [19] Handbook of Chemistry and Physics, 36th Edition, CRC Publications, Cleveland OH (1954)
- [20] M. A. Green, et al, "Forced Two-Phase Helium Cooling of Large Superconducting Magnets," Advances in Cryogenic Engineering **25**, p 420, Plenum Press, New York, (1979)